Crypto Tension

Description: We start off this week with a fabulous Picture of the Week and, for the first time in this podcast's 12-year history, our first Quote of the Week. Then we'll be discussing the chilling effects of arresting ethical hackers, the upcoming neutrality debate congressional hearing, something troubling I encountered at McAfee.com, an entirely new IoT nightmare you couldn't have seen coming and just won't believe, the long-awaited Adobe Flash end-of-life schedule, welcome performance news for Firefox users, the FCC allocates new sensor spectrum for self-driving cars, three bits of follow-up errata, a bit of miscellany, and then Crypto Tension - a careful look at the presently ongoing controversy surrounding the deliberate provisioning of passive eavesdropping decryption being seriously considered for inclusion in the forthcoming TLS v1.3 standard.

High quality (64 kbps) mp3 audio file URL: http://media.GRC.com/sn/SN-621.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-621-lq.mp3

SHOW TEASE: It's time for Security Now!. Steve Gibson is here. We're going to hear about a sad story of a young security expert who got busted in Budapest for finding all the flaws in the new subway system. We'll also hear about a big fight going on right now in the Internet Engineering Task Force, and Steve will explain it all. Plus iRobot that wants to sell your floor plans to the highest bidder. What could that mean? It's all coming up next on Security Now!.

Leo Laporte: This is Security Now! with Steve Gibson, Episode 621, recorded Tuesday, July 25th, 2017: Crypto Tension.

It's time for Security Now!, the show where we cover your latest security updates. Of course this is the post-Black Hat, pre-DEF CON show, so you want to pay close attention. Steve Gibson is here, our security guru. Hi, Steve.

Steve Gibson: Leo, great to be with you again, as always.

Leo: Good to see you, good to see you.

Steve: This interesting time of the summer when we have our annual major Las Vegas events, all the hackers getting together. Normally there are press releases sort of happening along the way, and then we do deeper coverage after. For whatever reason,

there doesn't seem to be a lot of upfront press coverage, so I imagine that we'll be talking about what was revealed, probably next week.

Leo: One thing that I think Greg Ferro, who covers it, was talking about - maybe it was Roberto Baldwin, who is on his way to DEF CON. He doesn't go to Black Hat. Black Hat's more corporate; DEF CON's more fun. He said that because the bug bounties these days are so steep, people just don't hold onto security exploits in the way they used to; right?

Steve: Yes, exactly. So, yes, you can prance around onstage and have some glamour, or you could have a quarter million dollars. Gee.

Leo: Uh-huh, uh-huh.

Steve: What do I choose? But we do have a packed and interesting podcast. We'll start off this week with a fabulous Picture of the Week that you already saw and got a big kick out of. Then for the first time in this podcast's 12-year history our first Quote of the Week, which started off being just - I was going to cover it; and I thought, okay, I'm going to pull this out separately because this is just too fun. Then we're going to discuss the chilling effects of arresting ethical hackers, something that happened in Hungary; the upcoming neutrality debate congressional hearing; something troubling I encountered yesterday when I went to McAfee.com; an entirely new IoT nightmare we couldn't have seen coming and won't believe; the long-awaited Adobe Flash end-of-life schedule, thank goodness; welcome performance news for Firefox users; the FCC allocating some new spectrum for self-driving cars; three bits of follow-up errata; a bit of miscellany; and then arguably the main topic of this podcast.

So unlike the recent ones, where we just haven't had - there's been so much news and no one thing really stood out? This one does. I titled the podcast "Crypto Tension" because we're going to take a look at the end of the podcast, a careful look at the presently ongoing controversy, and I can't believe I'm even saying this, surrounding the deliberate provisioning of passive eavesdropping decryption being seriously considered for inclusion in the forthcoming TLS v1.3 specification. So, yikes.

Leo: Yep, mm-hmm.

Steve: But first, a word from one of our several loving podcast sponsors.

Leo: I wanted to mention, this is a terrible thing to do, but you liked "The Martian," right, Andy Weir's wonderful book and later movie?

Steve: Oh, loved it, yes, yes.

Leo: Weir's been working on a new book. We knew this. We had him on "Triangulation." We're getting him on "Triangulation" again in, I can't remember,

soon. And I just got a copy of his new book "Artemis." And I was all excited until I saw, "On Sale November 14th." So I don't want - I won't tease you. And I will be reading this.

Steve: And "The Martian" was one of those where I knew the movie was happening, so I deliberately read the book first.

Leo: Right.

Steve: And of course the book is always better.

Leo: It really was, yeah, in this case.

Steve: Yeah, I mean, there's so much extra detail that a book is able to have that you just sort of have to do in broad strokes or broad brush with a movie.

Leo: That's always the case.

Steve: Yeah.

Leo: Plus I've always held, especially with science fiction, your mind is a much better set builder than any real Hollywood set builder. And so these are much more vivid when you read them. Now, the success of the "The Martian" movie was such that he's already sold this. Movie rights have been acquired by 20th Century Fox, even though the book won't be out for months. It's about - and actually it will be a great movie. I don't want to give away too much. I will read to you from just the back of the book.

Steve: Yeah, yeah, yeah.

Leo: "Welcome to Artemis, the first and only city on the moon. It's a tourist destination and an economic miracle. It's a testament to our species' ingenuity, resourcefulness, and strength. If you've got the money, you can stay at the Ritz-Carlton Artemis, which is on the Aldrin bubble. If you're a little tighter, you could stay in the Conrad" - poor Pete Conrad, he got the cheap bubble. Anyway, it sounds awesome. I think that that's not a spoiler.

Steve: It does sound like a fabulous movie.

Leo: What a premise; right? What a premise.

Steve: Yeah.

Leo: I won't read more, although there is more on the back of the book. But I was so excited. I just opened this, so I'm very excited. I don't think Mark Watney is in this one, though. But I'm sure Andy Weir's incredible sense of humor is.

Steve: I don't think Mark ever wants to leave the Earth again.

Leo: Mark is staying put, staying put. All right. Enough of that.

Steve: Okay. So our very first in 12 years Quote of the Week. Todd Westby, who is the CEO of Wisconsin-based tech company Three Square Market, was explaining to a reporter from ABC News...

Leo: Oh, dear.

Steve: ...about how 50 of the company's 80 employees had agreed to be "chipped."

Leo: Oh, I saw this.

Steve: As is the slang term, have an RFID tag the size of a grain of rice implanted in that little webbing between their thumb and forefinger.

Leo: We were talking about this at breakfast, Lisa and I. She said, "Would you get chipped?" I said, "Well," I said, "it's not a GPS. It's probably RFID," which it was.

Steve: Okay, so, yeah. So I just love the way he's tried to explain this, like sort of why this is a good idea.

Leo: It's like chipping your pet. It's the same exact technology.

Steve: Well, yes. For our audience, though, he says to this reporter: "There's really nothing to hack in it because it is encrypted just like credit cards are."

Leo: Oh, please. What? What?

Steve: And then he says: "The chances of hacking into it are almost nonexistent..."

Leo: Oh, this guy doesn't know what he's talking about.

Steve: I know, "...because it's not connected to the Internet," he said. "The only way for somebody to get connectivity to it is to basically chop off your hand."

Leo: Oh, god, please.

Steve: Oh, that makes me feel so much better. Those bullet points, that just sells the whole concept.

Leo: Well, the thing that strikes me is that's, I mean, it's just sending out a number; right? It's like a credit card number that's [crosstalk].

Steve: It is a transponder. It is a passive transponder.

Leo: So you put it near the reader, it'll send it a code.

Steve: You ping it with a little magnetic pulse that gives it the energy it needs in order to send back a fixed ID.

Leo: And it's a hard-coded number. That's my problem. It's like a fingerprint. I mean, you're not going to - it's not like a credit card, which you can change if somebody gets it. It's more like a fingerprint, which you're out of luck if somebody gets it; right?

Steve: Well, it was described as a splinter, so it can be removed.

Leo: I guess you could remove it and put it back.

Steve: If somebody regrets it. He also said that his wife, his young adult children, and others would be getting, you know, probably Rover, as well, will be getting the microchip next week.

Leo: Oh, lord.

Steve: And, you know, what could possibly go wrong? So again, I get it that people are split on this. There could be - and he's saying, oh, you'll be able to just, like, put your hand up against a reader at the front door to enter and so forth. I mean, I get that there's a convenience aspect. My problem is, if it in fact produces a fixed ID, that's, as we know, far less secure than a rotating ID. It can't...

Leo: I doubt that it does rotating. It's too small; right?

Steve: Well, it could. It can't be time based because there's no way for it to power a clock. But it could be like those - remember the early VeriSign eInk cards, where every time you pressed the button you got a different six digits. So there would be enough energy to increment a counter. And with a little supercapacitor to keep it alive, blah blah blah.

Anyway, the point is that there are ways to make it better. The problem with it being a fixed output is that then it's cloneable. The problem is you might tend to over-rely on the security of something which is really not very secure. This is not a secure technology. It's a convenience. But if you also just memorized a long passcode, that would be safer, and nobody would want to cut off your hand. Which is a serious upside to not being chipped. Anyway.

Leo: By the way, this is old technology. I remember Mexican diplomats who were subject to being kidnapped did this, not so much for credit cards, but just so that they would be identified when they got the remains.

Steve: Well, yeah. And of course we're seeing an evolution of the technology. The original RFID tags were like a large pill, I mean, like a really large capsule. And so it was much more inconvenient to have that somewhere. And they were, like, putting it in your forearm because you needed a bigger place. But that little webbing in between your thumb and first finger, that seems to be now the place to lodge...

Leo: I'm willing to bet a hundred bucks that it's exactly the same technology you do with Fido when you put a chip in his ear, that all it is is an ID number.

Steve: Yeah. I think you're absolutely right.

Leo: Yeah.

Steve: Yeah. And by the way, they cost $300 each. They're also not cheap.

Leo: It's profit center.

Steve: That's right. So speaking of something that is cloneable - we'll get to that at the end of this story. Budapest's Hungarian Public Transportation Authority - whose name I cannot pronounce, but the initials, thank goodness, are BKK - turned in an 18-year-old ethical hacker after he notified them of a trivial exploit to their recently put online ticketing website.

Leo: Oh, nice.

Steve: So, okay. It's just so wrong. So get a load of this. The 18 year old, who has asked that his name not be made public, was poking at a newly available mobile online ticketing system which had just been put online. So he just you know, wanted to see how

they did. And in other coverage of this that I read, the sudden appearance of this surprised people because this transportation agency had been trying to get e-ticketing going for years and had spent, I think it was something like 12 million equivalent U.S. dollars in apparently nothing so far. So suddenly this thing appears. Okay. He discovered, this 18 year old, that after bringing up the BKK's website, he could simply press F12 to open the browser's built-in developer tools, modify the page's form submission code to alter the ticket's price.

Leo: Oh, this guy was an elite hacker, obviously.

Steve: Oh, that's right. We've got to lock that guy up because you don't want to let him loose on society. And because there was absolutely no client- or server-side ticket price validation in place, the BKK system blindly accepted the visitor-provided ticket price.

Leo: Name your own price.

Steve: I mean, it was almost as if, I mean, it's the equivalent of putting a blank field on the page saying, "Fill in the price you would like to pay."

Leo: How much would you like to pay?

Steve: And it issued a valid ticket at that reduced price. That is, charged him - well, and so as a demo the young man says he purchased a ticket, normally priced at a U.S. equivalent $35, for just $0.20, to see if he could. Now, he didn't even know it was going to work. I mean, of course you wouldn't know looking at this. You would hope it wouldn't work. So he adjusts the form's statement of the ticket price from the equivalent of $35 to $0.20 as a simple proof of concept, and he never used the ticket in any way.

And of course, as I just said, when he initially made the trivial change to the web page, he didn't know if it was going to work. It was a trial. It was a test. Well, after responsibly notifying the transportation authority of his finding, so that they could address this glaring deficiency in their brand new, just-put-online system, and never using the valid ticket, shortly afterward he was awakened in the middle of the night and arrested by the police.

Leo: That's sad.

Steve: It is. BKK Management then boasted in a press conference about, quote, "catching the hacker" - yeah, it was tough, we read our email - and declaring their systems secure once again because of this...

Leo: Because we got the guy.

Steve: Yeah, we got him. He's behind bars. Bleeping Computer's reporting said that, since then, other security flaws - you can imagine, if it does this, what else must be

there. Other security flaws in BKK's system have since surfaced on Twitter. And I put in parens here, yeah, no kidding. I said I would hate to be BKK right now, if their publicly facing systems design is so slipshod. Can you imagine what else must be wrong there?

And get this. BKK has a $1 million annual contract with a local company, T-Systems, for the maintenance of its IT systems. So somebody's got a sweetheart deal there with the government. Okay. But we actually do have additional results. Since then, obviously, the news of this has drawn a lot of attention. We have learned that the system stores its passwords in cleartext and emails them in the clear, if you ask for a reminder. I forgot my password. Oh, here it is. Oh, thank you very much. That's what - I forgot it was Aunt Bessie's maiden name. Also, after logging in, visitors were able to get the data of other users, apparently simply through manipulating the page's URL.

Leo: Oh, lord.

Steve: So your account name is showing in the URL. And if you change it to somebody else, oh, look at that, now you're them.

Leo: Oh, my goodness.

Steve: Reports have claimed that it's possible to access other users' profiles, which include their full name, their physical address, and an ID number which is either their national ID, their driver's license, or passport.

Leo: And of course not easily changed. No, wow.

Steve: Exactly.

Leo: Wow. Wow.

Steve: And if you put in "shop.bkk.hu" into the browser URL, nothing happens because the site never implemented an HTTP to HTTPS redirect to bounce the browser's default assumption of HTTP, which by the way I've commented previously ought to be changed now. Browsers should try HTTPS or try both and choose the one that is secure, if they both work. But that's not happening yet.

So if somebody hears or reads that, oh, the new e-ticketing system is shop.bkk.hu, and they go to their browser and type it in, they don't get to the page because they have to do https://shop.bkk.hu explicitly because these people didn't bother to open port 80 and bounce the user over to 443. Also, tickets don't even work. They don't display properly on iPhone's Safari browser. So that's a problem. And, finally, someone determined, I don't know how this could possibly have happened that they figured this out, another case of master hacking, that the admin password was adminadmin.

Leo: No.

Steve: And logged in using that.

Leo: Oh, whoever did this was a sixth-grader. Wow. Wow.

Steve: Oh, and, finally, it doesn't even work. Getting back to the RFID tag problem, there's no tracking of ticket usage or cancellation of a ticket upon use. So tickets were 100% copyable, and there were reports you hit the home button and power to take a snapshot of the ticket being displayed on your phone, and then email it to other people. And...

Leo: Everybody can use it.

Steve: Yes, exactly.

Leo: One ticket serves all.

Steve: Some guys did an experiment, made a video showing the reuse of the same ticket through "ticket control," unquote, 10 out of 10 times, without being caught or raising any alarm, and the ticket was accepted. So the BKK representatives in their announcement talked about how the system was under continuous attack, they said, of which none were successful.

Leo: No.

Steve: So consequently there was no need to stop the system, and that everybody's data was safe. Nothing to see here. These are not the droids you're looking for. Move along.

Leo: Oh, wow, wow.

Steve: Wow. So we know that criminal cyber hacking is a very real thing today. Bad guys won't disclose the breaches they find. They'll exploit them to the limit, doing real damage over time to their victims. Compare that, contrast that to a teenager who verifies a problem and privately, quietly, and responsibly reports it to the affected company, thus allowing them, if these people had - well, I mean, now we realize it literally was the tip of the iceberg he found, but would allow them to fix the problem with no fanfare. You know, I get it. When you look at the laws, the laws are awful. The laws that exist now are clearly wrong. And it is a difficult problem because we know there is a wide gray area that separates the white hats and the black hats. And you and I have talked about instances, Leo, in the past where a hacker kind of does some things that are questionable.

Leo: Gray hat, yeah, yeah.

Steve: Yeah. And then kind of says, oh, but I really didn't mean it, or I was just checking, it's like...

Leo: I was pen testing, yeah, yeah.

Steve: Yeah, yeah, yeah. And it's like, well, were you really? So I get it that there's a gray area. But what we have currently are very overbroad laws which are still on the books and label anything that some authority in power doesn't like, basically. I mean, that's almost the way the law is written. If you don't like this being done to you, then it's illegal. But the problem is we need to provide, I mean, it's so crucial that an 18 year old somehow be able to not do damage to a potential victim, but help them, be allowed to help them. They're paying a million dollars to some ridiculous company who has allowed this to happen, and here's an 18 year old who, for free, says you might want to fix this because anybody can set their ticket price.

Leo: Everybody's embarrassed. They're just embarrassed, that's all.

Steve: Yeah, exactly. It is, you're right, it is a bureaucratic kneejerk reaction by people who don't understand any of the technology.

Leo: Right. They're just embarrassed.

Steve: And so we have to take the control of this out of their hands.

Leo: We hope cooler heads prevail, and this poor kid gets out of jail. Hungary is sliding badly into authoritarian government, like others in this world.

Steve: Which will remain nameless for the moment.

Leo: Yeah. But so I don't have the highest hopes. This is kind of a symptom of authoritarian government.

Steve: And in a sort of related case of, whoops, that's really not what we meant, and I heard you guys talking about this on Sunday, whether or not ISPs like it, Title II is still in effect. And Verizon was recently caught deliberately violating it.

Leo: Yeah, mm-hmm.

Steve: Now, they're claiming, after the fact, that they were...

Leo: Just a test. It was just a test.

Steve: Now, but they didn't announce it beforehand.

Leo: It was a test to see if we'd get caught.

Steve: Yes. They didn't say, oh, we'll be running some tests next week. In fact, it even took them a while to generate their corporate response to being outed. They said they were testing performance optimizations for video content on their network, whatever that means - actually, I think we know what that means, slow down video - and the effect was clear. Netflix, YouTube, and other video streaming services were being throttled.

This was not disclosed by Verizon until after it was discovered and became public. Users who were achieving 30Mb download on an LTE connection of non-video content were measuring a reduced Netflix data rate of a flat 10Mb. I mean, that's what Net Neutrality is trying to prevent. Which is to say that Verizon was clearly conscious of the source of the data and was limiting its speed on their network because they could.

Now, also people during this time did some experimenting. Using a VPN unthrottled the connection, since Verizon was then unable to peer into the traffic and did not know that it was coming from Netflix or YouTube or wherever, and thus could not throttle it based upon its source. As we covered last week, one of the major ISPs said that they wanted their 2015 Title II classification as a common carrier repealed - and I love this - and that, when that was done, they would honor Net Neutrality voluntarily.

Leo: Yeah, we'll do it, we'll do it. Just make sure we don't have to do it. That's the problem. Yeah, we don't want to have to do it.

Steve: Exactly. And here we appear to have Verizon breaking Net Neutrality while still under Title II, which makes that unlawful. So what possible hope is there, if ISPs are released from the legal obligation to treat all traffic equally? And we know that ISPs don't want users to use VPNs. You were also talking about that on Sunday. But if content and source-based traffic shaping, as it's called, I mean, that is the technical term, "traffic shaping," is applied to consumer data streams, an increasing use of VPNs is foreseeable. Consumers have a tool, and it's very difficult to block because now VPNs can run over port 80 and port 443 and look just like encrypted web traffic. So here's another interesting battle looming.

And speaking of Net Neutrality, I wanted to note that there was some nice coverage by Jon Brodkin in Ars Technica, and I'm so glad that this hearing is going to be on a Thursday and not on a Tuesday. The biggest websites and the biggest Internet service providers are being summoned to Congress to testify about Net Neutrality. The Chair of the House Energy and Commerce Committee, U.S. Representative Greg Walden, who's a Republican Representative of Oregon, said he's scheduling a full committee hearing titled "Ground Rules for the Internet Ecosystem" that will be set for Thursday, September 7th, so a ways away still. I'm sure we'll be reminding our listeners. This is one committee hearing I'm going to be watching.

This Tuesday morning, during an FCC oversight hearing, he said: "Today I'm sending formal invitations to the top executives of the leading technology companies including Facebook, Alphabet, Amazon, and Netflix, as well as broadband providers including Comcast, AT&T, Verizon, and Charter, inviting each of them to come and testify before our full Energy and Commerce Committee." Now, of course we know the question is -

and this is the point you brought up, which was exactly right, Leo, last week. The question is whether this is all just political theater, and whether the lobbyists may have already won this battle, and testimony is only being taken for face-saving purposes.

But also, as we said last week, for better or for worse, what we need here is clear and clean law, rather than the vicissitudes of presidential appointee mandates which are changing everything. And Walden did say, he said he wants Congress to step I, and said both ISPs and websites should weigh in first. He said: "It's time for Congress to legislate the rules of the Internet and stop the ping-pong game of regulations and litigation." Yay.

Leo: Yeah, I think that is the right answer. I really do.

Steve: Yes, yes. "Given the importance," he said, "of this public policy debate and the work we need to do as a committee, it is essential that we hear directly from the country's top Internet and edge provider leaders who frequently speak out publicly about rules of the Internet. It's time they came before us" - but haven't they before? - "and directly shared," he says, "their positions and answered our questions." That's why this is going to be must-see TV. "With more than a month's advance notice," he said, "I'm sure they can arrange their schedules to accommodate our invitations." And I'm sure they're going to want to because this is high stakes. I think this is important. And, as I said, thank goodness it's on a Thursday.

Okay. One more, and then we'll take a break, our second break. Or our second sponsor, our first break. My browser complained yesterday when I visited McAfee's website. So there was a little bit of news that I don't know if we covered a few months back. TechCrunch had a nice piece of reporting at the start of this past April. They wrote: "If you were on the Internet in a certain era, you remember McAfee." And of course none of us have forgotten John.

TechCrunch writes: "It was the defensive line between you and the rest of the Internet, reminding you with incessant pop-ups that you were not hacked, not quite yet, but only if you renewed your subscription right away. Intel bought the firewall company in 2010" - and I was thinking, wow, I didn't realize it was that far back, seven years ago - "for an eye-popping $7.68 billion and billed it as Intel Security, and the name McAfee became more closely associated with the company's founder, a man who retired to Belize only to be accused of his neighbor's murder." And then, as a little bit of trivia, Johnny Depp will apparently be playing John McAfee in an upcoming film. That should be interesting.

Leo: Perfect casting, yeah.

Steve: Yeah. But then TechCrunch writes: "But things didn't work out with Intel" - and then they said, parens - "(or Belize, either, for that matter). So the unit formerly known as Intel Security will be" - and this is them writing in April - "McAfee once again. Today, Intel is officially inking a deal that will spin McAfee out, with the asset management firm TPG taking a 51% stake" - so a majority share - "in the company for 4.2 billion," so a little more than it was purchased for in terms of percentage back in 2010 by Intel. So it didn't jump up in value much. Intel will retain the balancing 49% share.

Now, here's where McAfee stands. McAfee currently secures two thirds of the world's 2,000 largest companies and grew its revenue 11% in the first half of last year, in the first half of 2016. So it's a going concern. Okay. Yesterday I go to McAfee.com, right, the

enterprise security firm. And I'm greeted with an across-the-page fixed-position floating bar of text stating that, quote, "Your browser is blocking some features of this website." This is on McAfee.com. "Please follow the instructions at" - and then there's a URL - "support.heateor.com/browser-blocking-social-features to unblock these."

So I'm thinking, what? And I, like, double check. Am I at McAfee? And, yeah. So, like, okay. I did a little digging and made sure this Heateor deal was legitimate, and I went there, clicked the link that was on the McAfee.com site. So at the top of the page at Heateor, and this is a tongue-twister, it reads: "Sassy Social Share, Super Socializer WordPress."

Leo: Oh, lord.

Steve: From McAfee. Headline: "Why Is My Browser Blocking Social Features of the Webpage?" This is dated March 17 of this year, 2017, so it's recent. "Your browser," I'm reading, "might be blocking Social Features of the webpage you are facing issues with" - okay, the enterprise security firm securing two thirds of the top 2,000 enterprises - "related to loading social content," says this. And I put in parens here, "Oh, no!"

"If you are using Mozilla Firefox browser [uh-huh, yes] and it has Tracking Protection feature enabled [of course it does] you may have issues," this page reads, "in getting content loaded from social media websites such as Facebook, Twitter, et cetera. These features include social share counts, social avatars, social comments, and social login." Yes, all those social things we so desperately need from our McAfee Enterprise Security provider.

And then the page says: "To get the social content unblocked, you need to disable Tracking Protection of Firefox by following the steps mentioned below." And then they explain how I go to the location bar and enter about:config, and then "This might void your warranty" page may appear. Say "I accept the risk..."

Leo: Oh, my god.

Steve: ...to continue. Then search for "trackingprotection." Then double-click on privacy.trackingprotection.enabled to set it to FALSE.

Leo: Oh, please.

Steve: Anyway, enough said. Who knows? This may have come from an ad. I mean, it's probably some widget somewhere that some person, I won't use a more descriptive adjective, put there. Okay.

Leo: Is it still there?

Steve: Try going to McAfee. Well, you'd have to have Firefox...

Leo: I have Firefox. But I should disable tracking, huh? Where is it?

Steve: Yup. And then it came up. And it eventually went away because I was trying to make a screenshot of it, and I had already captured the URL, so I was able to see what it was. But, yeah, it was...

Leo: Is that the settings, or do I do that in...

Steve: About:config.

Leo: About:config, okay.

Steve: And then put into the search bar "trackingprotection." And that should eliminate all - it should whittle down all of the possible things...

Leo: I am getting "This might void your warranty."

Steve: Yeah, there you go.

Leo: But that's coming from Firefox. That's them saying, you know, if you're messing with this, you could be - let's see. Track. Safe browsing. Tracking protection is not enabled, so I'm going to enable it; right? Now it's TRUE.

Steve: Right.

Leo: All right, now let's go - let's have some fun. Let's go to...

Steve: McAfee.com.

Leo: ...McAfee.com. Oh, yeah, tracking protection. No, that's good. Next, okay, that's just Firefox saying what's going on on the page. That's fine, got it, okay. Ah, you know what, they must have taken that off.

Steve: Scroll down a little bit. See if you see something not moving on your page.

Leo: I think they must have gotten a few complaints. They do have, you know, social sharing widgets at the bottom, like I would share, oh, I was just on the McAfee page.

Steve: Oh, you know what it was? I followed a link there. This is not - the home page is not the page I was at.

Leo: Oh, okay.

Steve: And unfortunately I don't know what the page was. It didn't occur to me that it...

Leo: So I'll browse around. I'll see if I can get it to fire off again, yeah. Yeah, they want you to share this on Facebook. What they need is a like. They probably want to have a Facebook Like button. That's probably what it is. You know, thumbs up, I like this page. I like it.

Steve: Yeah.

Leo: All right.

Steve: So I found the page. I realized Firefox had the whole history. And there's a shortcut link that - McAfee's own sort of equivalent of bit.ly. And so you'd be able to type it in. It's not coming back up for me, though. But it did disappear after it had been there for a while. So https://mcafee.ly.

Leo: Ah, McAfeely.

Steve: Forward slash, and then numeral 2, lowercase u, uppercase A, uppercase A, 8, uppercase S, uppercase X.

Leo: Okay. This is millions of Android devices hit with copycat malware.

Steve: Yup, that's the page. And are you - and so you do not see it.

Leo: Oh, I do see it.

Steve: Oh, there it is.

Leo: Your browser is blocking some features of this fine website. Please follow the instructions. I don't blame you for being baffled because who wants to go to H-E-A-T-E-O-R, Heateor.com?

Steve: Yes. And it's darkened the page. And if you scroll, it's fixed. And so it's just sitting there like...

Leo: Oh, this is terrible.

Steve: This is McAfee.com. Two thirds of the top 2,000 enterprise companies.

Leo: So this is, I would guess, you see the social sharing block on here, it has LinkedIn, Facebook, Twitter, Google+, RSS. I would guess that they want you to be able to share. Oh, and they also support Facebook comments. So maybe that's the issue. I don't know. Wow.

Steve: Yeah.

Leo: I think that is so irresponsible to say, you know, why would you want to block tracking?

Steve: Well, except - yes, exactly. They're taking you to a link telling you to void your Firefox warranty and turn off tracking protection. It's like, okay.

Leo: I don't want that stuff, yeah.

Steve: That's good security. Thank you, McAfee. Maybe John is around more than we think.

Leo: Well, you know, I've recommended against this company for years. It's not - it's dubious value.

Steve: Okay. So here's a headline that just takes your breath away. "Roomba maker..."

Leo: Oh, yeah.

Steve: "...preparing to sell maps of your home to advertisers."

Leo: A new form of revenue.

Steve: You just can't make this up. Yes, it's funny, too, because Mark Thompson, a good friend of both of ours, was testing the roaming features of various floor vacs, and he actually showed me the - I don't know how he even did it. But he was somehow, like, tracking the pattern of motion of various models and demonstrating, like he had screenshots of the entire path that a Roomba took through his home versus something else. And it was surprisingly competent in its navigation. I was very impressed with what I saw.

But so yesterday iRobot CEO Colin Angle announced plans to sell maps of users' homes to advertisers. In 2015, iRobot started selling Roomba models capable of mapping homes, so the vacuums would know where they should go and stop bumping into furniture and other things. Until now, these maps have been kept and used only internally on the device to aid its navigation and its understanding of the environment. But iRobot realized there was a monetization possibility there and now plans to upload the maps of its customers' homes to its servers, from where they will be sold to online advertisers, apparently not endpoint advertisers but companies like Amazon, Apple, and Google. As I understand it, and as Bleeping Computer reports, the primary buyers aren't regular ad companies, but makers of smart home voice assistance like the Amazon device, the Apple device, and the Google device.

The idea is, and it's a little unclear to me, that these companies would buy this data, that is, the mapping data of the homes their devices are in, because you can imagine there's probably a relatively good correlation between people who have those and also have Roombas roaming around. They would buy the data and combine it with the telemetry they already obtain from their devices to build more sophisticated user profiles that they in turn - like maybe square footage. Bigger houses are more valuable and so forth, more TVs in them, who knows. In turn, they can sell, down the road, the plan is, to classic advertising companies or offer...

Leo: This is made up. Wait a minute. Slow down. This is what Bleeping Computer is assuming. This is not what's going on.

Steve: Correct, correct.

Leo: So I'll give you an example. Wouldn't it be nice if your stereo knew how to shape your acoustics? This is, by the way, what Apple's Home Hub will do, based on furniture positioning in your house. It's not, I mean, really, is there a big privacy violation knowing where your divan is in the house? And, by the way, Roomba says we don't do it currently. We're thinking, we're seeing if there's interest in it, and we will of course seek proactive approval before we do it.

Steve: Right.

Leo: But I don't - I think that, you know, you can think of all sorts of, I guess, all sorts, a few nefarious uses for this. But I can think of some very useful uses for this. Or even just what size should a sofa be? If we're going to make sofas, what size do people want sofas to be? Are we making them too big? Can we get around the sofa?

Steve: And again, I always come down, you know, you and I do differ on the privacy aspects. But my position has always been constant. As long as the user is informed, I have no problem with this.

Leo: Right, right.

Steve: If it's proactive, and they are giving permission for the floor plan of their home to

be exported and out of their control, then I don't have a problem with that.

Leo: I had the same reaction when I read it, like what? But I'm trying to think of, yeah, of course you would - and as we know, companies often bury this stuff in Terms of Service. And we've got to pay attention to this and make sure the Roomba does ask for permission. And, by the way, the company's started selling military robots. So who knows? This might be an urban assault feature. I don't know. But as long as you're informed and you can opt out of it - I think the default is to opt out. But again, that's what we prefer.

Steve: Well, no. They want - this is the problem, Leo. There is pressure, now that they turn it into profit, there is pressure for them to get it from their users.

Leo: Right, sure.

Steve: So this tilts the balance in a way that's uncomfortable. And we don't know how the maps are going to be protected.

Leo: Right, right.

Steve: Maps are going to be now - the maps of your home will be out of your control in the cloud, subject to search warrants, so that the government can say we want to get a map of this person's home.

Leo: But what would they do with that?

Steve: I'm just saying this is what happens. All of your other hosts are where I am on this, Leo. I've listened to them say...

Leo: No, I know they are. But I think there's a certain amount of techno panic, like oh, they can do this. Oh, my god. But I would like people also to think about, well, okay, I understand just kind of the default kneejerk response is nothing about me can be discerned by anybody. But what's the, I mean, really, what are they going to do with this that's going to harm you?

Steve: I don't disagree. My whole deal is make sure people know because it'll be very interesting to see, when this happens, how iRobot handles it because they just need to do it responsibly. None of us...

Leo: Well, frankly, they've announced that they're going to do this. So clearly they're not doing it in the shadows. They say they're going to do it; right?

Steve: Correct. Correct.

Leo: They didn't hide it.

Steve: Well, okay, except the CEO announcing it is very different from their consumers knowing it.

Leo: Right.

Steve: Those are two different things.

Leo: No, no, they need to - I agree. They need to very clearly specify this, yeah.

Steve: Right. Yeah, so anyway, I love it as content for this podcast because what we're seeing is we're seeing this creeping technology Internet-connectedness of virtually our entire, what was our private life, with devices that are listening to us and looking at us and recommending what we should wear and adjusting themselves to the size of our room, thus knowing the size of our room, and what's the contents of our refrigerator.

Again, I'm not saying that it's clear, but this is sort of this creeping exportation of or exfiltration of every manner of information. And it does get interesting when you start aggregating it. I mean, apparently that's what these home technology providers want to do, Amazon and Google and Apple. They want to, ooh, look, here's some additional interesting information we will now be able to purchase, and we'll know who it is because they all have relationships with the users whose homes their devices are in. So they'll associate the person's home floor plan with their devices. And again...

Leo: I can see why Amazon might want that for more effective array mics, things like that.

Steve: And actually you came up with some good use cases, like, oh, look. Couches are generally too large for people's homes; okay?

Leo: Right. I'm challenged to think of a nefarious use case. I mean, you came up with one, which is, well, then we know you have a fancy home. But it's just a map of the floor plan. It's not - I guess they know what your square footage is, but they can go to Zillow and find that out.

Steve: Very good point. Very good point.

Leo: I don't know - yeah. You know, in general I'm not too worried about stuff like that.

Steve: All I want to do is just kind of keep an eye on it. From my standpoint, I mean, I have split levels, and I don't have anything crawling - and lots of little rugs. So it's like a

death trap for a Roomba here. So I'm not being mapped, and I don't have any of these things listening to me. I don't need them. It's not a thing that I do. But just, you know, interesting evolution of our environment. And again, the idea that something that began mapping a few years ago, so that it could do a better job of mapping, and on that definition there's been some definition drift. Now it's, oh, look, we just realized we could upload these floor plans to our servers and sell them. So, okay.

The end of life for Flash has actually been scheduled, finally. And 2020 cannot come soon enough. Google has a blog. Mary Jo covered this in her column on ZDNet. And of course I refuse to run Flash anywhere. It's one thing to require it for video playback. I get that. Although, as we know, it's been possible to play video with pure HTML on a browser for so long now that there's really no excuse for requiring it. My site has been playing video Flash-free for years.

But I'm also confounded by non-video sites that have Flash helpers of one sort or another where, for example, my browser complains that a page is apparently attempting to run Flash for non-video reasons. I, of course, just say no and never wonder any further because I don't want Flash to run because we know the most likely scenario, and we've covered this on this podcast for years, is that an embedded advertisement is the culprit, and accepting third-party advertisers' Flash script is the last thing you want to do. Flash ads have been, and we've covered them for years, the primary vector of malvertising infection.

And wouldn't you know it. While putting this show together, I go to the Hacker News, which is TheHackerNews.com, a great site. They have super nice coverage, and their banner is The Hacker News: Security in a Serious Way. And my browser drops down a little bar saying "Allow http://thehackernews.com to run Adobe Flash." And it's, what, you know? So something there, and I don't know if it's their widget or a third-party widget that they've linked to, but my browser is saying this site wants to run Flash. How do you feel about that? And again, it's not providing any value to me, so of course I say no.

Okay. So this morning, this Tuesday morning, Google's Chromium team blog headline was, in homage to what I loved that Douglas Adams called the "fourth book in the 'Hitchhiker's Guide to the Galaxy' trilogy," Google's Chromium blog was "So long, and thanks for all the Flash." They said: "This morning, Adobe announced their plans" - oh, it must have been yesterday morning - "to end support for Flash late in 2020," so three years from now. "For Flash developers this will mean transitioning to HTML, as Chrome will increasingly require explicit permission from users to run Flash content" - now, that's from now until then - "until support is removed completely at the end of 2020."

This is Chromium's blog. "HTML is faster, safer, more power-efficient than Flash, and works across desktop and mobile. Three years ago," they said, "over 80% of Chrome's daily desktop users visited sites containing Flash. Today, only 17% of users visit sites with Flash." So in three years there's been a drop from 80% to below 20, down to 17. And they said, "And we're continuing to see a downward trend as sites move to HTML." They wrote: "We strongly encourage sites that still rely on Flash to make the move to HTML as there will be an increasing number of restrictions on Flash leading up to the end of support." So as Google always does when they do a campaign of this sort - we saw this with the SHA-1 certs; we've seen this with them working against the irresponsibility of some certificate authorities in the past - they do this in stages.

Mary Jo wrote that Adobe finally has drawn a line in the sand, noting that Flash will no longer be supported after 2020. Microsoft officials - so this is Mary Jo reporting on Microsoft - said that they'd do their part to wind down Flash support in the company's

Internet and Edge browsers, so that Flash support will be entirely removed from Windows by the end of 2020, as well.

She said: "Flash in Edge already is only click-to-run," she said, "as of the Windows 10 Creators Update. Today Microsoft posted," she wrote, "its timeline and plan for getting rid of Flash over the next three years." So through the end of 2017 and into 2018, Microsoft Edge will continue to ask users for permission to run Flash on most sites the first time the site is visited and will remember the user's preference on subsequent visits. So that sounds like a proper tradeoff, not to overly harass people who really do still want to run it and need to. IE will continue to allow Flash with no special permissions required during this time.

In mid to late 2018, so around this time next year, they write, "we will update Microsoft Edge to require permission for Flash to be run each session." So they'll take away the sticky memory of your previous decision, and IE will still continue to allow Flash for all sites throughout 2018. Around this same time in two years, mid to late 2019, they write, "we will disable Flash by default in both Edge and Internet Explorer. Users will be able to reenable Flash in both browsers. When reenabled, Microsoft Edge will continue to require approval for Flash on a site-by-site basis."

By the end of 2020, finally, in sync with Adobe's final cutting off all additional support, they say, "we will remove the ability, completely remove the ability to run Adobe Flash in Edge and IE across all supported versions of Microsoft Windows. Users will no longer have any ability to enable and run Flash." And then Mary Jo concluded, saying that Google, Mozilla - so also Firefox - and Apple are also committing to dropping Flash support by 2020 and probably do something similar.

So yay, you know, HTML can do everything now that Flash was needed for. And there was arguably a good need for it. Back in the day, it was the way we first did video and first did lots of things that HTML could not do. Now it's just built-in native. And it's become much more of a security concern than it is a benefit. And as Google said, it takes up space, consumes power, blah blah blah. And of course I guess that must be Apple on the Mac because iOS has always been Flash-free from day one, yeah.

Leo: Right. No, it's Safari, yeah.

Steve: And I just - I love this. A number of our listeners, a bunch of our listeners, I think, are still using - have not let go of Firefox, as indeed I have not. There's good news I just wanted to briefly share to those of us who are sticking with it and who enjoy organizing with tabs, as we know I do. I've actually - I used to be keeping about 212 tabs open, from literally as far back as when I had paused the work on the next version of SpinRite in order to switch to SQRL. They were all still there, all the tabs I had open at the time, because that's just sort of how I manage my stuff. You should see my desktop. Or, no, actually you shouldn't.

Anyway, June's NetMarketShare shows Firefox commanding now only a 12% share of the browser universe. And as we know, Chrome has been eating everyone's lunch. Chrome is nearly 60%. They're at 59.5% of the market share. And I guess a lot of that might be, what, Android and Chromebooks. But a lot of Windows and Mac users are also Chrome users. So it is a popular add-on browser. And certainly, whenever you go to Google, if you hit the home page, they're trying to get you to use Chrome, if you're not there with a Chrome browser. I see that all the time. So they're trying to push their browser. And we know it's a good browser. It's got a great security model. My problem with it is it is

seriously resource intensive. I mean, it gobbles memory. So I just - I can't run it.

Leo: You find that Firefox is more efficient now? It used to be a memory hog.

Steve: Yes. And so here's what's happened. If you're interested, Leo, the link under this, the TechRadar link has some graphs showing the evolution of Firefox in terms of performance and memory over time. Mozilla's so-called "Quantum Flow" project is bearing fruit in the next release number 55 of Firefox. When loaded down, get this, with a massive test case of - and this even makes me tremble.

Leo: I did read this.

Steve: 1,691 open tabs. I can't even imagine it getting off the ground. 1,691 open tabs. Okay. But it actually does. The current version of Firefox number 54 required four minutes to start and consumed 2 gig of system memory, which actually is immense, seems low to me, but anyway. By comparison, the next release, Firefox 55 with Quantum Flow technology, started the same daunting test set of tabs, not in four minutes, but in 15 seconds, and consumed less than half a gig of RAM. So dramatically faster and dramatically reduced. I think there's a second graph somewhere I saw. It must have been on some other...

Leo: I thought I saw that, too, yeah.

Steve: Yeah, some other reporting of this. So anyway, that's all I had to say. I get it that we're now the minority browser. I love my Firefox. I love my tabs on the side. If Chrome - and I'm a little constrained because I'm still in a 32-bit OS, so I'm capped at 3 gigs. I just, as we know, the next machine I build is either 64 gig or 128 gig of RAM, so it will be Chrome friendly.

Leo: I thought you built your last machine.

Steve: Oh, I did, but I'm not using it yet. It'll take me too long to switch over, and I've got to get...

Leo: You actually built that machine, but you're still not using it?

Steve: Yeah. Oh, because it'll take a chunk of time to get it set up and configured.

Leo: It's not as a server. It's your desktop; right?

Steve: Yeah, yeah. But remember, it doesn't run any 16-bit code, Win7 doesn't. And I've got legacy stuff. Like Brief, I can't use Brief under Win7.

Leo: No, yeah, you can't.

Steve: Which is my code editor.

Leo: You need to just set up a virtual machine with DOS 5 on it.

Steve: And now we're starting to understand why it will take me time, which is why I haven't made the switch yet. Okay, so...

Leo: Actually, be an interesting project to have everything run in a container isolated within that system. I think you could do that.

Steve: Well, and Qubes, the Qubes OS.

Leo: Qubes OS, oh, yeah, that's how it works.

Steve: Very much like that.

Leo: But I bet you you could have Windows 10 running and have Docker containers, or maybe Hyper-V containers, but I think Docker containers for all the different apps you want, including DOS.

Steve: Yes. And that system is built.

Leo: It'd be all isolated, sandboxed.

Steve: That system is built to be a VM host. That's why I have, you know, someone says "64 gig?" It's like, yes, so that I can have many large gig OSes running at the same time. So that's the plan. And that will be my last system. But I'm not using it yet.

Leo: But is it assembled?

Steve: It's right next to me, yeah. It's over there. Remember I showed pictures, and I talked about...

Leo: It's just dark, huh?

Steve: Yeah, I just, I haven't, I mean, it hasn't been updated in like, I don't know, 18 months, ever since I - I just, you know...

Leo: Now, are you going to - you're not going to put Windows 10 on it, then.

Steve: Oh, it's got 7 already set up.

Leo: Seven, okay.

Steve: I will never go to Windows 10. And I will clarify a statement I made in the errata, I think it is, in a minute, about Windows 10. But first, the FCC has just approved a sizable new chunk of radar spectrum for use by vehicular environment sensing radar.

Leo: Ooh.

Steve: Yes. This will enable the use of reduced cost and increased precision sensors in our next generation of autos. As we know, many consumer vehicles already use radar. Even if they're not going to drive themselves, they use it for collision avoidance, automatic lane-keeping and so forth. But right now, the Washington Post writes, vehicular radar - oh, and my god, the pun of the year. I take my hat off to Brian for this. I'll get to it in a second.

But he writes: "Vehicular radar is divided into a couple of different chunks of radio spectrum. Last Thursday, the Federal Communications Commission voted to consolidate these chunks and added more to allocate additional bandwidth to vehicular radar. FCC Commissioner Clyburn said: 'While we enthusiastically harness new technology that will ultimately propel us to a driverless future, we must maintain our focus on safety, and radar applications play an important role.'"

So then Brian Fung wrote, again, this is just such a good - this is the best pun. He said: "Thursday's decision" - that's last Thursday's decision - "by the FCC lets vehicle radar take advantage of the spectrum ranging from 76 GHz to 81, reflecting an addition of four extra gigahertz." Oh. Because of course that would be an awkward phrase if you didn't realize, of course, that radar is a reflective technology. Radar works by reflection. So anyway, Brian...

Leo: Reflecting it at - you think he meant that?

Steve: Oh, he had to. "Reflecting an addition of four extra..."

Leo: Extremely nerdy.

Steve: It's wonderful. You wouldn't say "reflecting an addition of four extra gigahertz" unless you knew what you were saying.

Leo: Wow.

Steve: And the fact that radar is reflective, anyway, it's like, oh, bravo. Anyway, what we have had before is 24 GHz.

Leo: I think his math is off, though. But other than that - is that five?

Steve: Oh, you're right, it is.

Leo: Okay.

Steve: That would be five.

Leo: He got the pun, but he didn't get the math.

Steve: Yeah. Or, well, actually four extra, I'm not sure what the range at 24 gig is. Because it didn't show a range.

Leo: Oh, so they might have lost, yeah, so they probably lost some, yeah.

Steve: Yeah. But what's important is the band. What I wrote is, I said, I'll note that the increase in radar frequency - because we're going from 24 gig up to 76-81. The increase in radar frequency is significant. From an engineering standpoint the move from 24 to 76-81 significantly increases the resolving power of the radar, and the higher frequency means smaller and more efficient devices. We're all familiar by analogy with audio speakers, famously known as woofers and tweeters. We know that to produce low audio frequencies in free air requires a large diameter speaker cone, commonly known as a woofer, but that higher frequencies can be efficiently generated with smaller speakers.

This analogy holds at microwave frequencies, where this factor-of-three upward jump in frequency means a factor-of-three decrease in wavelength, which allows for many more array sensors within the same area. So this is a win both for radar assist and for future autonomous driving. So yay to the FCC for formally making this new chunk of spectrum, very high-frequency radar, available. It will give our cars better eyes.

Leo: Good, good.

Steve: I did want to remind our listeners, I got a ton of appreciative feedback from our listeners about the Humble Book Bundle that we talked about last week. I just wanted to give another reminder in case somebody said ooh, and then like forgot about it. I gave it this week's bit.ly link to help anyone get there: bit.ly/sn-621. So this episode number, 621, all lowercase, sn-621. That will take you there. Just under six days remaining. It was six days at 11:00 a.m. this morning Pacific time, so by the time you're hearing the podcast next week it's over. So I just did want to let everyone know. And a listener built a Humble Bundle Downloader that makes it easy to grab all of the bundle assets and other materials that are associated. It's at GitHub. I imagine you could just google or search "humblebundle-downloader" at GitHub, and I also have the link in the show notes

for anyone who's interested. So thanks for that.

Leo: Yes, indeed.

Steve: So first piece of errata.

Leo: Before you do that, I might have an errata. You never showed the Picture of the Week.

Steve: Oh.

Leo: And it's so good, I didn't want it to slide by.

Steve: Oh, I'm glad you reminded me. Yes, yes, yes, thank you.

Leo: Not to mention timely, yeah.

Steve: Thank you. Let's do that, and then we'll take our last break.

Leo: Okay. Okay.

Steve: Yes. Oh, thank you, Leo. You loved it, and I loved it, but nobody else got to hear about it.

Leo: It was just the two of us, yeah.

Steve: That's right. So this is so perfect. This is a picture, several people sent it to me, and I got a couple of them really compressed, but I got a good one. This is a placard standing up on the little sponge mat on the counter at the UPS Store in Las Vegas.

Leo: This week, yeah.

Steve: This week, yes. And this placard, it's got The UPS Store, and then their logo printed in color. It says: "Due to the DEF CON Hacking Convention, we will be accepting email print jobs with attachments only. We will not accept USB prints or any links." And then, down below, "We apologize for the inconvenience."

Leo: I'm not sure that's actually better, but okay.

Steve: Well, and they had to have been bitten in the past.

Leo: Maybe that's it, yeah.

Steve: It had to have happened. And so they're like, okay.

Leo: Oh, those hackers. Oh.

Steve: I mean, maybe they're proactive, and they're just really on the ball, and they realize the danger. Or maybe some scary-looking people came in.

Leo: So funny. So funny.

Steve: Oh, wonderful, wonderful.

Leo: Don't fall for those spurious links. You might be pulled to a McAfee page or something worse.

Steve: Okay. So it turns out a report that everybody covered, including we here, was completely specious.

Leo: Oh, no.

Steve: Yes. No audio device called the police when that guy was threatening his wife with a gun. Wired magazine, fortunately, followed up on what turned out to be a widely reported, entirely erroneous story about an unnamed home audio device which we'll remember was originally believed to be a Google Home unit, and then people believed it was the Amazon Echo, autonomously responding to a very loud and fraught domestic dispute by phoning the police. Except it never happened.

Leo: Yeah. I was a little suspicious, I think you'll remember, yeah.

Steve: Yes. I mean, but there wasn't, well, and the reason everyone believed it is it's what the authorities said.

Leo: Right.

Steve: So Wired wrote: "Despite what you may have heard, an Amazon Echo did not call the police earlier this week, when it heard a husband threatening his wife with a gun in New Mexico. On Monday, news reports took Bernalillo County authorities' version of those events credulously, heralding the home assistant as a hero. The alleged act also

raised an important question: Do you really want to" - and they're not coming down on either side of this. "Do you really want to live in a world where Alexa listens to your conversations and calls the cops if she thinks" - ooh, I said the word, sorry.

Leo: That's okay. She's calling the cops right now.

Steve: "If she thinks things are getting out of hand. The good news is that you don't live in that world. Amazon's device" - and I've replaced her name with that word. "Amazon's device can't, and did not, call 911. Google Home can't do it, either. No," writes Wired, "voice-assistant device on the market can. That doesn't invalidate the core question, though, especially as Amazon's Echo, Google Home, and their offshoots increasingly gain abilities and become more integral to everyday life. How intrusive do you want to let these devices be?" asked Wired rhetorically. "Should they be able to call the police? Maybe not even just when specially prompted, but because they may have heard, for instance, a gunshot."

Okay, so "The Bernalillo County incident almost certainly had nothing to do with Amazon's Echo, but it presents an opportunity," Wired writes, "to think about issues and abilities that will become real sooner than you might think." And so anyway, "The Sheriff's Department reported specifically that, when a man drew a gun on his wife in a home where an Amazon Echo was placed, he said to her, 'Did you call the sheriffs?' and the Echo misinterpreted that, they said, as a command to call the sheriffs, who then showed up at the front door. The authorities later clarified that someone in the house could be heard in the 911 recording yelling the A-word, 'Call 911.'"

Leo: Ah. They'd already called 911.

Steve: Well, so they're now saying somebody was heard instructing the Amazon Echo to do that. Then Wired says: "That could not have happened, either. Amazon's Echo requires first a wake word to activate." As we all know, the default is the A-word, but you can also customize it to Echo, Amazon, or Computer. "And while they can make calls, an Alexa-powered device can only call another Alexa-powered device. Not only that, but it can only call other Alexa devices" - god, I keep saying the word, sorry, because I'm reading the text that I wrote - "can only call other Echoes..."

Leo: Just say Echo, yeah.

Steve: "...that have enabled calling and have been added to your contact list. Most importantly, these exchanges don't take place over the public switched telephone network, the worldwide network as we know that allows wireless and land phones to actually make calls. In other words, the sheriffs would have needed an Echo device of their own for that to ever work, one that the couple in the domestic dispute had in their contact list," which frankly seems really unlikely in this case, in the case of that household. "Later, the police said that the Echo device was used in combination with some kind of home phone or cellular," I mean, I don't think these guys know one end up from the other.

Leo: They don't know what's going on.

Steve: Yeah. And so they're, like, they keep changing their story.

Leo: "Oh, I think it did."

Steve: So we don't know what happened, but I did want to correct the record. It wasn't the case.

Leo: The Roomba did it.

Steve: It might have been sweeping in the room and looked up. Oh, my goodness. Okay. So some people were confused...

Leo: Oh, there's my front door. Hold on, let me check this.

Steve: Some people were confused about the example I used of how ISPs could acquire certificate interception capability when I used Google's ability to mint their own certificates as an example. So I just wanted to correct the record and make sure that people did not think that I was suggesting that Google was minting certificates for domains they don't control. Absolutely not. That was never my intention. And I'm sure they never have and never would. They're the good guys. But they are minting their own certificates, and that means they could. And so I wanted to make that distinction very clear.

So that an ISP, if it were decided that that was what we wanted, and this was the technological architecture that I wanted to explain, an ISP could be given on-the-fly certificate-signing capabilities for domains they don't control, namely the ones their customers are visiting. So I just wanted to make a very clear line there, that I used Google as an example of somebody whose certificate is signed by a CA where the certificate that was signed has CA authority, and an ISP could get such a certificate. They don't have them now, we assume and hope. But it's possible.

So I only meant that to explain the technology and that it wouldn't require, as I was suggesting previously, that all of our devices accept a root cert from our ISPs. There's a much worse and more powerful, but more practical from a technology standpoint, solution. Also, I thought that iOS's update 10.3.2, which we talked about last week, fixed that Broadpwn bug. Remember we talked about how Google's patch fixed it in Android. That was the baseband firmware problem found in Broadcom's WiFi chip that allowed anybody using basically a QoS-style packet, just you don't even have to log onto the access point, you just had to be in contact with its radio, and your vulnerable device could be owned, that is, could have code executed on it.

Anyway, it turns out that there was a fix related to the baseband firmware in 10.3.2, but not that one. It was the subsequent 10.3.3 iOS release, which just came out, which fixed that and a whole bunch of other problems. So I wanted to correct the record.

Two bits of crazy miscellany. One is that I've seen the Season 2 trailer of the next, well, Season 2 of Netflix's "Stranger Things." So the good news is we're going to get one.

Leo: Yeah, yeah.

Steve: We knew we were. And the boys are back. And this is one of those where you want them to create a lot of content quickly so they don't grow up because they're at the perfect age right now. And I don't want them to be teenagers because it just won't work as well as it does with them being in elementary school. So we are going to get a second season. Not till October, but the first trailer is out.

And we talked at length last week about the various sites which are blocking our password managers and just cutting, copying, and pasting into forms. There is an extension available for both Firefox and Chrome, which I can't give the name of on the podcast because the second word is the F-bomb.

Leo: Oh.

Steve: So Don't "F" with Paste. And so expand the "F" to the word we all know it stands for: Don't F**k with Paste. That is an extension available for Firefox and Chrome and maybe elsewhere. I know of those two because I've had users who are listeners tell me they're using it for both of those different browsers. So apparently there is a way to overcome that annoyance with that extension, which un-effs what the website has done, which is handy.

Leo: Don't "F" with Paste.

Steve: So, and I have one of those real heartwarming SpinRite data recovery stories. We've been talking a lot about the technology of SpinRite recently and what it does and data recovery and SMART data and how SMART data is only useful when the drive's under load, and numbers of ECC and all that kind of stuff. Steven Almas sent me just one of those where I'm just like, I'm so glad it does these sorts of things. He wrote: "My neighbor, who is a single mother, asked for my IT help. She had an external hard drive with 60,000 photos of her family and children, which was her only copy." And you know where this story is going.

Leo: Oh, oh.

Steve: "The hard drive did not mount on any computer, and she was very upset." 60,000 photos. He said: "I suggested that we purchase SpinRite and try that out. Since I work in IT, I took a dedicated machine and ran SpinRite on her external hard drive, and SpinRite was able to recover all but 13 of the 60,000 photos." He said: "She now has a comprehensive backup solution in place, and we are forever grateful to Steve Gibson and SpinRite." And he signed off saying, "Thank you from another Steve." So Steve, thanks for sharing it. And again, that's just so cool, 60,000 family photos of her and her children that would have been lost.

Leo: Wow.

Steve: And, yes, SpinRite's $89. But you have it for life. I'm never letting it die. As soon as SQRL is put to bed, I will be right back to it. And as I have said, everyone who has a version of 6 will be able to play with the next release before its formal release, and there will be no charge for the updates that I do in succession to 6, creating .1, .2, and .3 and so forth. And then my plan is to take it down completely and rewrite it from scratch, adding a whole ton of new features to it. And again, of course, even people who have SpinRite 1 are able to upgrade to 6 today. So I will be honoring upgrades moving forward. So anyway, Steve, thanks for sharing that. And, boy, that's so cool, to be able to recover, what would it be, 59,987 photos out of 60,000, all but 13.

Leo: And the 13 he couldn't recover, that's just because those were the damaged sectors, or some [crosstalk] photo lived on it; right?

Steve: Well, yeah. What this says is, and we see this with external drives, they get kicked around.

Leo: Yeah.

Steve: They fall off the table. They get knocked over if they're standing up on edge. They tend to get beaten up. So my guess is that there was probably growing, accruing damage over time, and the drive was struggling. And again, because there was no maintenance being performed, I mean, had she already had SpinRite, this would have - there would have been zero loss because, if you run SpinRite occasionally, it proactively, as we know, it maintains this and prevents it from happening.

So I'm glad you brought that up, Leo, because the fact that there was not even complete recovery says there was a lot of damage. And so SpinRite did everything it could and no doubt struggled, even on those final 13 photos, but just finally said they're just gone. The data is not here. The head was bouncing around on top of them when it fell off the table or something. And so, I mean, I really see that with removable drives because people don't appreciate them. When it's in a computer, even in a laptop, the inertia of the laptop tends to protect the drive, which is shock-mounted inside, to varying degrees. There isn't much room for much shock mount. And so laptops are a problem, too. But the problem is, if it's a little removable drive, it's just - it's under real G-shock threat.

Leo: Mm-hmm.

Steve: Okay. Two quick pieces of closing the loop. Actually, one of them is quick, and one is not. But we have time. Ned Griffin and several others said @SGgrc - so it was a tweet. He said: "Steve, did I hear you say 'my Win 10 machine'? Never thought I would hear you are using a flying turd OS." Which of course I did famously characterize Windows 10 as at some point in the future, I mean, at some time in the past. And so, yes, I have a Windows 10 machine, and I have been talking about Creators Update and the changes that are coming and so forth. And I recognize I have my own bias, but that can't be allowed to influence unduly my reporting, for the purpose of this podcast, what's

going on with Windows 10.

So I've got a Win 10 machine. I updated just because I wanted to experiment and see that "I forgot my password" link that we were talking about coming to the login lock page. And of course I'm also - I need to make sure that SQRL works seamlessly and correctly under Windows 10. So it does. But I only know that because I've been testing it all along the way. So, yes, I'm sitting in front of XP. But I do have all versions of Windows around me because I'm a developer, and I need them.

Also, a couple of our listeners, and this is an ongoing flux in my Twitter feed from our listeners, is when they see a site with obviously bad security practice, they just shoot me a tweet saying, hey, just thought you'd be interested in knowing. So, for example, while I was pulling this show together, I ran across someone who sent me a note saying that the gov.uk website allows a maximum of 12 characters and no symbols. And then someone else sent a login page screenshot that showed a range. I think it's maybe like eight to 15.

And so I just wanted to take a moment to put this into perspective. This is sort of an opportunity. This is well known to a lot of our listeners, but I know we have people who haven't been listening forever. So I wanted to remind everyone that, while these sorts of limitations pose a definite concern, because they give us cause for concern about the underlying security practices of the site's design and technology, in and of themselves they are insufficient evidence either way.

That's one of the things we need to remember. We're just seeing part of the front of a much deeper process. The part we're seeing does not give us confidence, but it also doesn't tell the whole story. For example, a site could allow a super long password, no limit, type until you're exhausted, but then store it directly in plaintext, which the user submits. So if their database were to leak, and we have covered account database leakage for years here, it would be game over for all of a site's users, despite the unlimited password complexity offered upfront. So there's an example of it really doesn't matter how much space they give you.

And on the flipside, a site could restrict their input password to, say, just eight characters, then pair that with a per-account large random nonce, which we know prevents a single attack against all of the site's passwords. If then that nonce is combined with what the user provided and is run through, as would be possible, a monstrous memory-hard, acceleration-resistant, like five-second-long hash process to make every brute force guess, whether offline or online, impossibly slow and costly. So that would turn eight characters into impossible to practically break.

So again, we're only seeing the front, and the back could render the most complex password unsafe, or the least complex password arguably more safe, by making its guessing extremely difficult. So anyway, I just sort of thought, because I'm constantly seeing people saying this, it's like, yes, it's certainly the case that any site telling you to fix the limit to a certain thing, that's a worry because that suggests, strongly suggests, that they have allocated that much space in their database. And we all know that, if you are hashing a password, any password as you should, then a hash by its nature turns a variable-length thing into a fixed-length result so that it doesn't matter how long the input password is. The result is always the same length, and so it can be stored in a fixed-length database field. So definite cause for concern, but not in and of itself clear evidence one way or the other.

Okay. Now, this is just sad. The title of the podcast I named "Crypto Tension." And this is the TLS v1.3 Explicit Wiretap Controversy. And before I got myself read into this enough,

I was a little concerned because at first blush it looked like our friend Matthew Green, the well-known cryptographer, was endorsing what he's calling "Data Center Use of Static Diffie-Hellman in TLS 1.3." That's the title of the IETF draft of his analysis of what's called "Static Diffie-Hellman." I'll explain, of course, all about what this is. So in his paper, and I've got the link in the show notes for anyone who wants the whole thing because I'm just snipping out a couple of the best parts.

The abstract reads - and this is Matthew Green wrote this. He's the sole author on this: "Unlike earlier versions of TLS, current drafts of TLS 1.3 have instead adopted ephemeral-mode Diffie-Hellman and elliptic-curve Diffie-Hellman as the primary cryptographic key exchange mechanism used in TLS." Okay, now, what he's saying there is that, unlike the earlier ones, which did allow some of these ciphers, but were primarily using RSA, TLS 1.3 is going to be primarily using elliptic-curve Diffie Hellman and ephemeral Diffie-Hellman for its key exchange mechanism. So he writes: "This document describes an optional configuration for TLS servers that allows for the use of a static, meaning an unchanging, Diffie-Hellman secret for all TLS connections made to the server. Passive monitoring of TLS connections can be enabled by installing a corresponding copy of this key in each monitoring device.

"While ephemeral elliptic-curve Diffie-Hellman," ephemeral EC Diffie-Hellman, he writes, "is in nearly all ways an improvement over the TLS RSA handshake, it has a limitation in certain enterprise settings, specifically the use of ephemeral PFS" - and that's, of course, the abbreviation for Perfect Forward Secrecy that we've talked about many times - "cipher suites is not compatible with enterprise network monitoring tools such as intrusion detection systems that must passively monitor Intranet TLS connections made to endpoints under the enterprise's control. This includes TLS connections made from enterprise load balancers at the edge of the enterprise network to internal enterprise TLS servers. It does not include TLS connections traveling over the external Internet. Such monitoring is ubiquitous and indispensable in some industries, and loss of this capability may slow adoption of TLS 1.3."

Okay. So let's understand. He, as a cryptographer, he's not taking a position. He's very involved in the TLS v1.3 working group. There has been a huge controversy, and I'll get to that in a second, about this, about this request which is coming from the enterprise side ostensibly, to break a set of important security guarantees which were going to be incorporated into TLS 1.3 for the purpose of monitoring, for the purpose of being able to passively decrypt TLS-encrypted traffic.

So I'm going to skip a bunch and just say that, in this document, the fourth point, and I've skipped over one, two, and three, of security considerations, he says: "We now consider the security implications of the change described above." So again, he's just saying, okay, if this is going to happen, let's make sure we don't do something wrong. Like, let's understand what this means.

He says: "The shift from fully-ephemeral elliptic-curve Diffie-Hellman to partially static Diffie-Hellman affects the security properties offered by the TLS 1.3 handshake by eliminating the perfect forward secrecy property provided by the server. If a server is compromised, and the private key is stolen" - and I'll just add, or if it is given - "then an attacker who observes" - that is, passively eavesdrops - "any TLS handshake, even one that occurred prior to the compromise" - or the gift or the servicing of a search warrant - "will be able to recover traffic encryption keys and be able" - that is, in the past - "and will be able to decrypt traffic.

"Thus the modification described in Section 4," Matthew writes, "represents a deliberate weakening of some security properties. Implementers who choose to include this

capability should carefully consider the risks to their infrastructure of using a handshake without perfect forward secrecy. Static secret keys," he says, "should be rotated regularly." Yeah. And of course they should be generated randomly per handshake. That's the way it is now. We're talking about breaking that on purpose.

So a great summary was produced by Stephen Checkoway, who's an assistant professor in the Department of Computer Science with the University of Illinois at Chicago. And I'm going to read this fast. But I can't skip it because he just - it's beautifully done. And with the background I just gave you, this will make more sense: "As the TLS 1.3 standardization process (hopefully)," he says in parens because it's been nine years, "comes to a close, there has been some drama on the TLS Working Group mailing list and at the recent IETF 99 meeting in Prague regarding the use of TLS 1.3 in enterprise networks. This is a surprisingly contentious and important topic that I suspect many people who don't follow protocol development closely may have missed." Thus you're getting it on this podcast.

"Transport Layer Security," he writes, "is, without exaggeration, the most important security protocol in use on the Internet today. It is the successor protocol to the older SSL protocol and is used to cryptographically protect a wide variety of Internet communications including online banking, a significant fraction of email traffic, more than half of all web browsing" - and of course we know that number is going up fast - "and an ever-increasing amount of normal Internet activity.

"TLS is standardized by the Internet Engineering Task Force (IETF) which is organized into a set of working groups. Each working group has a charter which describes its mission. The TLS Working Group is currently charged with designing the fourth iteration of the TLS protocol, TLS 1.3. This multi-year process takes place primarily on the TLS mailing list, as well as in regular, in-person meetings. The 99th IETF meeting just concluded. Sounds pretty dry. What's the drama about?" he says.

"Much of the work is pretty dry and technical. One of the working group's goals for the TLS 1.3 is to produce a more secure protocol than prior versions, which have had a series of subtle problems." As we know. We've covered them. "To that end, the working group has removed a number of cryptographic options that reduced the security. This removes options like cipher suites, sets of cryptographic algorithms that work together to secure the traffic that do not provide forward secrecy.

"To quote Wikipedia, 'A public-key system has the property of forward secrecy if it generates one random secret per session to complete a key agreement without using a deterministic algorithm. This means that the compromise of one message cannot compromise others as well, and there is no one secret value whose acquisition could compromise multiple messages.'" Perfectly phrased. "Forward secrecy also generally requires the session key to be destroyed once the session ends to prevent an adversary from decrypting traffic afterwards. So not only can you not decrypt traffic afterward, you cannot decrypt traffic that occurred beforehand." They're talking about killing that. Calm down, Steve.

"Forward secrecy is a very desirable property," he writes - Stephen writes. Was it Stephen Checkoway? Yeah. He writes: "Forward secrecy is a very desirable property in a cryptosystem. As I recall, when removing the non-forward-secret cipher suites was proposed on the mailing list for TLS 1.3, there was broad consensus. At some point, late into the TLS 1.3 design process, some enterprise network operators began to realize that this would reduce their ability to inspect traffic in order," they claim or they say, they state, "to troubleshoot problems within their networks and started asking the TLS working group to restore some of the removed cipher suites or provide some other

mechanism to support their internal network requirements."

He says: "The most recently proposed mechanism uses what's called 'static Diffie-Hellman'" - and this of course is what Matthew's first paper that I talked about was looking carefully at - "and works by," he writes, "reusing encryption keys. Interestingly, a form of this is used today as a minor optimization and isn't technically forbidden by TLS 1.3. Initially, the working group refused to consider any proposal which would hurt or remove forward secrecy. Recently, as the TLS 1.3 standardization effort has begun to draw to a close, the enterprise network operators have become more vocal. On the mailing list and at in-person meetings, three viewpoints have emerged. The debate between those with conflicting points of view has been vigorous," he says, "and in terms of the sheer number of words written, quite lengthy. What exactly do the enterprise folks want?

"In a nutshell, these enterprise operators want the ability to decrypt the traffic that is inside their own networks. Let's call this the enterprise viewpoint. Now keep in mind any network traffic that is inside their network was either, A, generated from inside their network, in which case the enterprise's own computers created the plaintext in the first place; or, B, the traffic was sent from the Internet to one of the enterprise's computers. In either case they already have the ability to do whatever they want with the plaintext," meaning because they're either endpoint where the decryption occurs naturally, "including storing all of it and examining it at will.

"If they already have access to the plaintext, why do they need changes to TLS 1.3 to enable them to get plaintext? The question is key to the whole debate. The enterprise viewpoint holds that operators need to be able to decrypt traffic from packet captures from various vantage points within the network. For example, they would like to decrypt traffic before and after a load balancer, web server, or database server in order to pinpoint which part of the network infrastructure is causing problems. On the mailing list and in person, they have been adamant that decryption from packet capture rather than, say, endpoint logging is the only way they can perform this sort of network debugging at the scale they need, given the fragility of what appear to be mind-bogglingly complex network architectures."

He says: "It seems pretty reasonable to support this use case. What's the problem with accommodating their request? After all, this will only be for use inside their own networks. On the one hand, this is reasonable and is completely supported today using TLS 1.2. Indeed, one of the suggestions has been for network operators to continue using TLS 1.2 inside their networks if they need this capability. On the other hand, there's no technical way to confine proposals to enable decryption to a particular network or data center." Right. You know, once the data's encrypted, that's the way it stays until it reaches its destination. So you couldn't use anywhere TLS 1.3 if you used 1.2 anywhere. That is, on a given connection.

"There are two major concerns raised by those opposed to breaking or degrading forward secrecy. Let's call this the forward-secret viewpoint. One concern raised by those with the forward-secret viewpoint is that proposals such as the static Diffie-Hellman approach mentioned above will enable wiretapping which would violate the IETF's Policy on Wiretapping. Although that may be true - and this is hotly contested - some other technical mechanisms have been proposed which would make such wiretapping externally visible.

"The second concern is both more subtle and, I think, more compelling. TLS, and SSL before it, has a history of supporting weak cryptography, and this support has come back to bite us several times." And of course on this podcast it's an often-encountered theme

and problem. "The best of example of this is the export cipher suites. These used cryptographically weak algorithms, but were at one point in time the only cipher suites that could be legally exported from the U.S. Two decades after the use of export cipher suites should have ended, researchers showed how to abuse support for these deprecated algorithms in modern TLS libraries to perform man-in-the-middle TLS connections.

"The forward-secret viewpoint holds that the TLS Working Group should not standardize any weaker form of TLS; and, if this makes some network operators' jobs harder, tough. That's two viewpoints, enterprise and forward-secret. What's the third? Let's call the third viewpoint the "pragmatic" viewpoint. This viewpoint holds that, whether or not enterprise network operators really need the decryption capability, some of them really want it; and, since they really want it, they're going to do something to get it. It's strictly better for the mechanism to be designed in public" - thus Matthew's contribution - "following normal IETF procedures, than to be cobbled together by people whose focus is on operations and not necessarily on security. It's worth noting that at least one of the authors of the static Diffie-Hellman proposal mentioned above firmly holds the pragmatic viewpoint." And he didn't name him, but I wouldn't be surprised if that was Matthew Green.

"Which viewpoint," he asks and finishes, "is correct? Before I say which viewpoint I think makes the strongest case, I want to point out that I'm sympathetic to all three. The network operator's job is just not an easy one, or so I assume. It's definitely outside," he writes, "my particular area of expertise. If they say they need plaintext in order to do their job, I don't think I'm in a position to contradict them. The pragmatic viewpoint is quite compelling. All else being equal, I'd much rather have the IETF design a standard mechanism to support the network operators' needs than have a hodgepodge of homegrown, difficult to use, non-interoperable, and potentially insecure solutions.

"But as they say, all else is rarely equal. The Internet is for end-users, not for network operators. The protocols we design today will, for better or worse, be in use for decades. End-users have been paying the price for our mistakes and past compromises on security. As protocol and implementation deficiencies necessitate new network hardware and software, the network operators have paid their own price.

"To rebut the enterprise and pragmatic viewpoints, I need not take a security-maximalist view. The sense of urgency from the operators and the pragmatists is, I believe, unwarranted. Yes, switching to TLS 1.3 will prevent operators from doing precisely what they're doing today; however, there is currently no need to switch. TLS 1.2 supports their use case; and TLS 1.2, when used correctly, is secure as far as we know. Of course, the network operators won't receive the benefits of mandatory forward secrecy, but that is precisely what they are asking to give up in TLS 1.3."

Finally: "Designing secure protocols is hard." Yeah. I've been working on one for a few years. "To date, our best efforts have not been as successful as we would like. In my view, the only option we have is to design the most secure protocols we can to achieve our stated objectives. We may still get it wrong, of course. My hope is that, in 20 years, we won't, once again, be dealing with security issues we know about today. Instead, I hope we'll be dealing with a whole new set of security issues."

So I finished on the dot of 4:00, and I'm glad I got that out because this is an important thing that is happening. We will keep a watch on this. I don't know how it's going to come down. The good news is browsers - what this means is that the key being sent by the server as part of the negotiation would not change from one connection to the next on the same server, which means browsers would be aware if this was being done to

their user and could raise a flag.

The other thing that's possible, because it certainly won't be the case that the earlier versions of TLS won't also be supported, is that browsers could choose not to advertise their support of TLS 1.3. That is, typically you're able to selectively disable these various protocols, either in the browser or in the OS. And so users could decide, uh, I'm not using this spyware TLS. I'm going to only say that I know about 1.2, which would be, in a sense, it would be a version downgrade, but unfortunately a security upgrade, where you'd be saying, nah, I don't want 1.3; 1.2 is just fine. I want my perfect forward secrecy cipher suites.

So anyway, really interesting that here at the very end, as they're getting ready to cross the finish line, with 1.2 being nine years old now, that suddenly these guys are saying, wait a minute, we need to spy on our own traffic. And it's like, okay. Well, we were going to fix it so no one could. Now we're not sure what to do.

Leo: Yeah, wow.

Steve: Interesting.

Leo: Well, how long do these meetings go on for?

Steve: Oh, well, most of it is in mailing lists.

Leo: Oh, okay. They are actually - they did just - because one of our correspondents on Sunday, as you know, was there, Greg Ferro.

Steve: Yep.

Leo: So I know they were having a physical meeting. But, yeah, of course the mailing list will continue to debate this.

Steve: Yeah, I was glad he corrected himself. He said that the URL was in the handshake. And of course I immediately, wait, what? No, it's not. But then he did say that there were some aspects of the [crosstalk].

Leo: There is some leak of data. Yeah, he did say "URL." But partial, and he said "partial," yeah.

Steve: It's the domain name. We already know that in the SNI, Server Name Indication extension to TLS, you declare what host you're asking for.

Leo: Interesting. So that's the one thing that's leaking. People know you've gone to

Google, but not what you're doing there.

Steve: Precisely.

Leo: Steve Gibson, he does it again, ladies and gentlemen. Thank you so much, Steve. Everybody should go right now to GRC.com and buy a copy of SpinRite, if you don't already have one.

Steve: Keeps me afloat.

Leo: And we want to keep him afloat, god knows, keep him in coffee. All you have to do is go to GRC.com. And while you're there, of course, there's plenty of other things to do, lots of freebies. Steve gives everything else away, including his work on SQRL, Perfect Paper Passwords, ShieldsUP!, Don't Shoot the Messenger, DCOMbobulator, I can go on and on and on.

Steve: Even the upgrades to SpinRite.

Leo: Yeah, those are free, too. And of course this show's absolutely free. In fact, Steve spends some of his money doing this show on transcripts, for one thing. He gets great transcripts written of each show, so you'll find those there, along with the audio, 64Kb audio. We have audio and video at our site, TWiT.tv/sn for Security Now!. And you will find it, of course, wherever you subscribe to podcasts. In fact, I encourage you to subscribe because this is one you want the whole set. You want every episode of Security Now!.

Let's see. Anything else? We do the show Wednesday, 1:30, I mean, sorry, Tuesday, 1:30 Pacific, 4:30 Eastern, 20:30 UTC, so you can watch live if you want. We stream it live. Just, you know, it's the behind-the-scenes version of the show. That's pre-transcript and all of that. What else? You can join us in the chatroom. There's always a lively conversation behind the scenes at irc.twit.tv. Thanks, Steve. Have a great week, and I'll see you next time on Security Now!.

Steve: Fantastic. We will be post-Vegas DEF CON and Black Hat, and I imagine we'll have some fun coverage of that.

Leo: You didn't talk about Fruitfly. That's going to be revealed tomorrow for the Mac.

Steve: Correct. I actually, well, Leo, we went two hours and 13 minutes.

Leo: I know, it was plenty, I know. And we did talk about it on MacBreak Weekly, and I'm sure there'll be more to talk about after their presentation.

Steve: And actually I had a whole bunch of tabs, and I saved them for next week because I just - I tried to figure out how long this was going to go, but I thought this discussion of TLS 1.2 was really important.

Leo: Yeah. All 1,691 tabs saved.

Steve: Right.

Leo: Now they'll load faster, too. Thank you, Steve.

Steve: See you next week, buddy. Bye.